



Supporting Analysis Users in U.S. ATLAS

Kaushik De

University of Texas At Arlington

U.S. ATLAS Tier 2/Tier 3 Workshop, UC


August 19, 2009

Overview




- U.S. ATLAS has excellent track record of supporting users
 - US is the most active cloud in ATLAS for user analysis
 - Analysis sites are in continuous and heavy use for >2 years
 - We have regularly scaled up resources to match user needs
 - But, we do not know if we are at the right scale (we expect user activity to grow significantly when data arrives)
 - We need to steer more users to run analysis at Tier 2's
 - We do not have much experience with supporting users at Tier 3
 - Planning guidance comes from previous talks
 - Tier 3 issues will be discussed tomorrow
 - In this talk, review current user analysis support and operations
 - Discuss plans for future

Analysis Usage Growing



Growing Tier 2 Activity Since May






Key Ingredients for Analysis Support



- What structures do we have in place for analysis support
 - Monitoring and support teams
 - Proactive monitoring and site support at U.S. Tier 1 and Tier 2's
 - Distributed Analysis Shift Team (DAST)
 - ATLAS Distributed Computing Operations Shift (ADCoS)
 - Facilities organization
 - U.S. ATLAS Integration Program
 - U.S. ATLAS Operations Program
 - Testing and validation exercises
 - Jamborees, stress tests, site certification, HammerCloud...
 - Software support
 - U.S. Panda team
 - Users running analysis, very active through hypernews

Site Monitoring Through Panda



US ✓	880	5552	08-19 03:47	546	0	0	0	4988	0	557	40	0 / 0	15569	5552	26% 8% 19%
----------------------	-----	------	-------------	---------------------	-------------------	-------------------	-------------------	----------------------	-------------------	---------------------	--------------------	-----------------------	-----------------------	----------------------	------------

US Sites	Job Nodes	Jobs	Latest	Pilot Nodes	defined	assigned	waiting	activated	sent	running	holding	transferring	finished	failed	tot	trf	other
Site Name	880	5552	08-19 03:47	546	0	0	0	4988	0	557	40	15569	5552	26%	8%	19%	
ANALY_AGLT2 ✓	0	0	offline	0	0	0	0	0	0	0	0	0 / 0	0	0			
ANALY_BNL_ATLAS_1 ✓	49	4403	08-19 02:08	92	0	0	0	0	0	6	0	0 / 0	7453	4403	37%	8%	29%
ANALY_BNL_LOCAL ✓	0	0		40	0	0	0	0	0	0	0	0 / 0	0	0			
ANALY_GLOW-ATLAS ✓	0	0	offline	0	0	0	0	0	0	0	0	0 / 0	0	0			
ANALY_HU_ATLAS_Tier2 ✓	0	0	offline	0	0	0	0	0	0	0	0	0 / 0	0	0			
ANALY_IllinoisHEP ✓	7	0	08-19 03:46	4	0	0	0	62	0	3	2	0 / 0	36	0	0%	0%	0%
ANALY_LONG_BNL_ATLAS ✓	120	285	08-19 03:47	97	0	0	0	0	0	29	3	0 / 0	1838	285	13%	8%	6%
ANALY_LONG_BNL_LOCAL ✓	0	0		2	0	0	0	0	0	0	0	0 / 0	0	0			
ANALY_MWT2 ✓	216	423	08-19 03:47	19	0	0	0	2442	0	447	0	0 / 0	388	423	52%	49%	3%
ANALY_NET2 ✓	109	8	08-19 03:47	40	0	0	0	182	0	23	13	0 / 0	810	8	1%	0%	1%
ANALY_OU_OCCHP_SWT2 ✓	6	8	08-19 00:10	14	0	0	0	0	0	0	0	0 / 0	0	8	100%	0%	100%
ANALY_SLAC ✓	211	421	08-19 03:47	216	0	0	0	2302	0	49	22	0 / 0	3048	421	12%	3%	9%
ANALY_SWT2_CPB ✓	162	4	08-19 03:41	20	0	0	0	0	0	0	0	0 / 0	1996	4	0%	0%	0%
ANALY_UTA ✓	0	0	offline	0	0	0	0	0	0	0	0	0 / 0	0	0			

Distributed Analysis Shift Team – DAST



- User analysis support is provided by the AtlasDAST (Atlas Distributed Analysis Shift Team) since September 29, 2008. Previously, user support was on a best effort basis provided by the Panda and Ganga software developers.
- Nurcan Ozturk (UTA) and Daniel van der Ster (CERN) are coordinating this effort.
- DAST organizes shifts currently in two time zones – US and CERN. One person from each zone is on shift for 7 hours a day covering between 9am-11pm CERN time, and 5 days a week.
- Members in the North American time zone:
 - Nurcan Ozturk (UTA)
 - Sergey Panitkin (BNL)
 - Alden Stradling (UTA)
 - Jakob Searcy (University of Oregon)

ADCoS (ADC Operations Shifts)

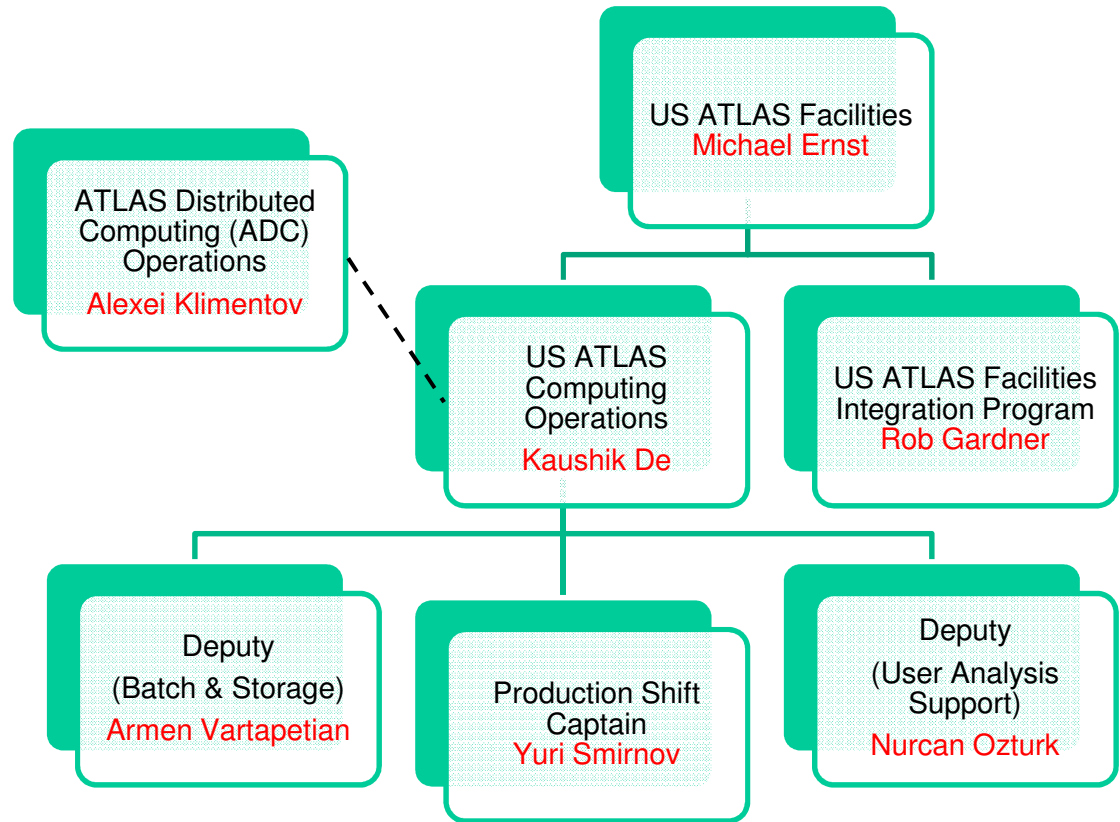


- ADCoS combined shifts started January 28th, 2008
 - Coordinated by K. De and Xavier Espinal (PIC/IFAE)
- ADCoS Goals
 - World-wide (distributed/remote) shifts
 - To monitor all ATLAS distributed computing resources
 - To provide Quality of Service (QoS) for all data processing
- Organization
 - Senior/Trainee: 2 day shifts, Expert: 7 day shifts
 - Three shift times (in CERN time zone):
 - ASIA/Pacific: 0h - 8h
 - EU-ME: 8h - 16h
 - Americas: 16h - 24h
- U.S. shift team
 - In operation long before ADCoS was started
 - Yuri Smirnov (captain), Mark Sosebee, Wensheng Deng, Barry Spurlock, Armen Vartapetian, Rupam Das

Facilities Organization



- See Michael Ernst's talk for overview
- Integration program covered in Rob Gardner's talk
- Operations activity started one year (+few months) ago



Data Distribution Issues




- Data distribution plan is in Jim Cochran's talk
 - How many copies of ESD/pDPD's do we need?
 - One copy at Tier 1, one copy among Tier 2's sufficient?
 - Do we have sufficient storage after first run?
 - How many copies of AOD/D1PD is optimal?
 - We are currently making 3-4 copies at US Tier 2's
 - How many copies of DnPD's generated by physics groups?
 - Not done uniformly or regularly yet
 - We expect these to be heavily used
 - Currently, we make 6 copies (Tier 1 + 5 Tier 2's)
 - Space may be a issue in the long run?
- Distribution of data on demand (RAW, ESD, HITS...)
 - Web interface from Panda monitoring page
 - Works well – but requests increasing rapidly
 - We do not have data expiration policy

Data Management



- Data movement through DQ2/Pandamover works well
- Cleaning up stored data is not easy
 - Data creation is well automated – production can fill up storage
 - Cleaning up is not so easy for a variety of reasons
 - Aborted and obsolete datasets now cleaned automatically in U.S. by ADC ops. (but physicists are reluctant to declare data obsolete)
 - Charles Waldman's scripts to check consistency and clean up
 - Armen/Hiro working on automated user data cleanup strategies
- Data placement not optimized
 - Production jobs have default storage either to tape or disk
 - Once produced, no mechanism to archive rarely used data to tape
 - Currently also missing simultaneous archiving to tape

Testing and Validation



- **Site certification and validation**
 - Through integration program (organized by Nurcan Ozturk)
 - Variety of test jobs, new ones added regularly
- **Jamborees**
 - Regularly organized at BNL, ANL and LBL
 - Leads to spike in real user analysis
 - See Jim Cochran's talk
- **Stress tests**
 - Many U.S. stress tests in the past
 - Moving towards coordinated tests on all clouds
 - HammerCloud helps with automation
- **STEP09**

Analysis Site Certification in US



Goal: certify sites to handle each type of analysis job in advance of the stress tests.

- Job preparations
- Site preparations and ANALY queue profiles
- Metrics

<http://www.usatlas.bnl.gov/twiki/bin/view/Admins/AnalysisSiteCertification.html>

Certifications by job type

Notation: completed work is in progress defer to next phase table to be updated

Job type	Athena release	input	T1	AGLT2	MWT2	NET2	SWT2	WT2
SusyValidation	14.5.0	AOD						
D3PD making with TopPhysTools	14.5.0	AOD						
TAG selection	14.5.0	TAG						
DB access job	14.5.1.4	AOD						
Data reprocessing	14.5.0.5	DPD						
ANLASC1	14.5.1	AOD						
ANLASC2	14.2.23	ESD						
ANLASC2 Jet sampling	14.5.1	ESD						
More jobs types to come ...								
HammerCloud Jobs								

200 jobs for each type, 95% success rate.



- Stress testing with HammerCloud
 - "Infinite" real analysis jobs to participating sites (all T2's + a few T1's)
 - Most (~all) of our STEP'09 results drawn from these tests
- Coordinated user activity
 - US and DE challenges
- Background (normal) user activity
 - Some sites were quite busy without HC's help
- HC workflows tested:
 - Classic Athena analyses:
 - various input data sizes: **ESD**, **AOD**, **DPD**
 - **TAG** for direct event access
 - Calibration & Alignment: **RAW** data and remote **DB access**
 - **Private MC** production
 - everything else (**ROOT**, **arbitrary executables...**)



HC: Global Summary




~1M jobs submitted -- 82.3% succeeded

Total 26.3B events processed


Mean Events/s = 7.7Hz

Mean CPU/Walltime = 0.39


Overall Efficiency



Overall CPU/Walltime



Overall Events/s



Global Efficiency: 82.3%

Global Rate: 28.6kHz



STEP09 - Cloud Summary



Cloud	# Jobs	# Successful	# Failed	Efficiency ▲	#files	#events	Hz	CPU/Wall
US	163004	153240	9764	0.940	465393	4169999722	6.0	33.9
IT	59163	52990	6173	0.896	311061	2798011153	6.5	32.6
NG	16730	14551	2179	0.870	20179	172708698	7.3	
FR	166427	144658	21769	0.869	557571	5050911395	9.3	44.5
ES	72562	62565	9997	0.862	236690	2150478621	10.0	44.3
TW	24178	19544	4634	0.808	86293	833817261	15.2	48.4
UK	181394	145222	36172	0.801	439084	3984651767	7.1	39.8
CA	41890	32306	9584	0.771	87117	757520054	4.2	34.4
DE	176076	135218	40858	0.768	557395	4991372555	7.9	40.8
NL	66632	37171	29461	0.558	154452	1352903529	7.9	35.0

In US sites have been busy with real user jobs.
- 426309 user analysis jobs during STEP09
- 173648 (~40%) from HammerCloud
- Remaining 60% from 198 other users
- 172k succeeded, 80k failed (of which 40-50k killed by user)

User support during STEP09 provided by DAST - Distributed Analysis Support Team:
- normal shift rotation w/o extra effort
- no extraordinary problems or complaints

Next Challenge



- During the summer we have generated massive MC data samples in the U.S.
 - `step09.00000011.jetStream_medcut.recon.AOD.a84/`
 - jet pt > 35 GeV, 97.69M events, 9.4pb⁻¹ integrated luminosity
 - total size 14900 GB, 9769 files
 - `step09.00000011.jetStream_lowcut.recon.AOD.a84/`
 - jet pt > 17 GeV, 27.49M events, 0.26 pb⁻¹ integrated luminosity
 - total size 3674 GB, 2749 files
 - Replicated to all US Tier2's
 - Nurcan Ozturk has tested and validated these samples
 - Additional 300M events generated (but not replicated to Tier 2's yet)
- Good sample for beginning students/postdocs

Scaling Up



- Tier 2 usage for analysis is still only ~20% of expected
 - Stress testing and HC has been very useful
 - But we need more real users to create 'chaos'
- Storage and data management issues
 - Required lot of attention during the past year
 - Slowly improving automation of system
 - Fundamental limits (to chaotic large scale data access) are still not fully explored – expect a lot of pain when data arrives
- Direct user support
 - Works well now, but funding is limited – we cannot suddenly add a lot more people to support users when data arrives (currently support model works well)
 - Need to continually automate (design goal of Panda from the beginning) all procedures – especially DDM

Conclusion



- Scaling up will require help from physicists
- We need to spend a larger fraction of our time on automation, instead of fighting fires with manual systems
- We have highly experienced operations team in place (many with >5 years experience with ATLAS grid)
- U.S. ATLAS is doing well – but there is still lot more work to be done before data arrives